

BFD Lab1 – eBGP

Important! This guide assumes that the AOS-CX ova has been installed and works in GNS3 or EVE-NG. Please refer to GNS3/EVE-NG initial setup labs if required.

<https://www.eve-ng.net/index.php/documentation/howtos/howto-add-aruba-cx-switch/>

At this time, EVE-NG does not support exporting/importing AOS-CX startup-config. The lab user should copy/paste the AOS-CX node configuration from the lab guide as described in the lab guide if required.

TABLE OF CONTENTS

Lab Objective.....	2
BFD reminder	2
Lab Overview.....	3
Lab Network Layout.....	4
Lab Tasks	4
Task 1 – Lab setup	4
Task 2 – Configure BGP routing	7
Step #1: Configure BGP	7
Step #2: Verify BGP sessions	8
Step #3: Verify BGP routes and ip routing table	10
Step #4: Verify IP connectivity between HostA and SW1 Loopback0	12
Task 3 – Test resiliency without BFD	12
Step #1: Start pings from HostA to SW1 L0 until interrupted	12
Step #2: Modify VLAN ID on interface 1/1/9 of SW2 and check traffic interruption	12
Step #3: restore VLAN 1105 on interface 1/1/9 of SW2.....	13
Step #4 (optional): test bgp fast-external-fallover	13
Task 4 – Configure BFD for the eBGP session established over the L2 path.....	13
Step #1: Configure BFD globally	14
Step #2: Configure BFD for eBGP session.....	14
Step #3: Verify BFD	14
Task 5 – Test resiliency with BFD.....	17
Step #1: Start pings from HostA to SW1 L0 until interrupted	17
Step #2: Modify VLAN ID on interface 1/1/9 of SW2 and check traffic interruption	18
Step #3: Check status of BFD session and BGP session. Check routes.	18
Appendix – Reference Configurations.....	21

Lab Objective

This lab will enable the reader to gain hands-on experience with configuring BFD on BGP and testing associated benefits on the network high-availability.

BFD reminder

- **Bidirectional Forwarding Detection** is a standard general-purpose fast failure detection mechanism.
- After a BFD session is established, if no BFD control packet is received from the peer within the negotiated BFD interval, BFD notifies a failure to the protocol consuming the BFD service, which then takes appropriate action (like tearing down BGP session, and consequently clearing the routes).
- On AOS-CX, BFD operates in asynchronous mode: BFD control packets are sent periodically to the peer. If the other device does not receive BFD control packet from the peer within the specified interval, it tears down the BFD session.
- Failure detection is unidirectional. The local node uses BFD control message to notify the peer that a BFD failure was detected.
- **Echo** is a function that consists in **looping back** the received echo message to the sender **without processing** it. If the sender does not receive its echo message back, BFD tears down the session. This task is accelerated in the ASIC, alleviating the CPU load for BFD control messages, as the rate of the control message can be lower. No delay of processing. It must be implemented in each direction. Echo function is not available on AOS-CX Simulator.
- IPv6 transport of BFD is supported on the AOS-CX Simulator.
- BFD should be used **only for non-direct point-to-point peering**, typically when a 3rd device is inserted between the peers. This 3rd device can be a L2 switch or a L1 optical switch not mirroring the state between physical ports.
- **Most of the time, BFD is not useful for BGP peering over a direct point-to-point L3 circuit:**
 - For **eBGP, fast-external failover** will ensure that if the peering circuit fails, all the routes learnt behind the next-hop attached to the failed circuit will immediately get withdrawn (in milliseconds).
 - For **iBGP, neighbor fall-over** command **without bfd** extension, will force route withdrawn based on the next-hop reachability status in OSPF. As the link connecting the peer fails, OSPF will immediately withdraw Loopback routes and notify BGP protocol which will do the same within milliseconds.
- When enabling BFD for BGP, the objective is to track the availability of the physical links being monitored and not the reachability of the loopback. Consequently, the loopback address must not be used as BGP peering IP address. It is not suitable for:
 - iBGP (traditionally peering over loopbacks)
 - eBGP multihop (TTL>1)

Lab Overview

This lab guide explains how to configure BFD (Bidirectional Forwarding Detection) on AOS-CX switch with eBGP.

As a prerequisite, please read the BFD section of the [ArubaOS-CX High Availability Guide](https://www.arubanetworks.com/techdocs/AOS-CX/10.07/HTML/5200-7854/Content/fir-int.htm) (<https://www.arubanetworks.com/techdocs/AOS-CX/10.07/HTML/5200-7854/Content/fir-int.htm>).

During this lab, you'll be able to:

- Configure eBGP over parallel paths and BGP weight as BGP criteria for best routes election
- Test traffic and path selection in nominal situation and also in failure condition.
- Configure BFD
- Check improvement on high-availability.

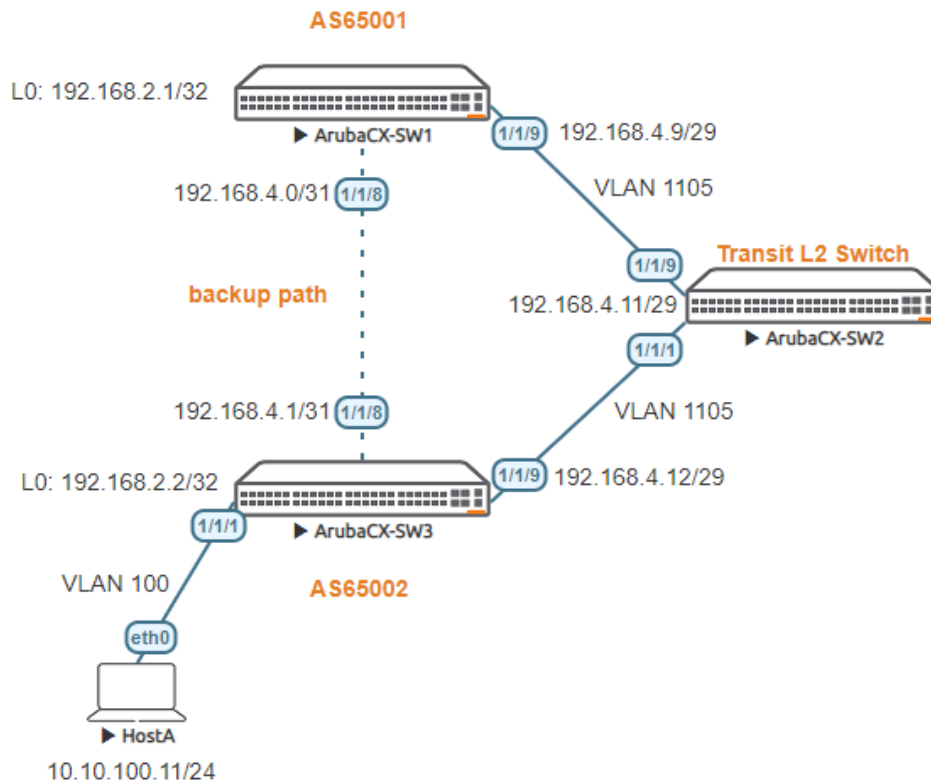
The minimum required AOS-CX Switch Simulator version for this lab is 10.6. It is recommended to use release 10.07.0010 or later.

This lab uses EVE-NG but GNS3 can be used as well.

Other BFD labs will be proposed like tracing and debugging BFD.

Lab Network Layout

Here is the proposed topology:



Lab Tasks

Task 1 – Lab setup

- In EVE-NG, import the .zip lab file containing the “unl” file.
All the connections between nodes are already set-up. Appropriate numbers of CPUs (2), RAM (4096 MB) and interfaces are already allocated.
- Check the connectivity as proposed above.
- Start all the devices (3 AOS-CX switches and 1 host)
- Open each switch console and log in with user “admin”.
The switches will ask to enter a new password. This new password can be an empty password for simplicity in this lab.
- Apply (copy/paste) the baseline configuration as proposed below

Baseline configuration proposal (for initial copy/paste):

SW1	SW2
hostname SW1 ! vlan 1 vlan 1105 description transit interco VLAN interface mgmt no shutdown ip dhcp	hostname SW2 ! vlan 1 vlan 1105 description transit interco VLAN interface mgmt no shutdown ip dhcp

```
interface 1/1/8
  no shutdown
  description link to SW3
  ip address 192.168.4.0/31
interface 1/1/9
  no shutdown
  description link to SW2
  no routing
  vlan access 1105
interface loopback 0
  ip address 192.168.2.1/32
interface vlan 1105
  ip address 192.168.4.9/29
```

SW3

```
hostname SW3
!
vlan 1,100
vlan 1105
  description transit interco VLAN
interface mgmt
  no shutdown
  ip dhcp
interface 1/1/1
  no shutdown
  description link to HostA
  no routing
  vlan access 100
interface 1/1/8
  no shutdown
  description to SW1
  ip address 192.168.4.1/31
interface 1/1/9
  no shutdown
  description to SW2
  no routing
  vlan access 1105
interface loopback 0
  ip address 192.168.2.2/32
interface vlan 100
  ip address 10.10.100.1/24
interface vlan 1105
  ip address 192.168.4.12/29
```

```
interface 1/1/1
  no shutdown
  description to SW3
  no routing
  vlan access 1105
interface 1/1/9
  no shutdown
  description to SW1
  no routing
  vlan access 1105
interface vlan 1105
  ip address 192.168.4.11/29
```

HostA

```
VPCS> ip 10.10.100.11/24 10.10.100.1
```

- Verify the connectivity through LLDP neighbor information as follows:

SW1

```
SW1# show lldp neighbor-info
```

```
LLDP Neighbor Information
=====
```

```
Total Neighbor Entries      : 2
Total Neighbor Entries Deleted : 0
Total Neighbor Entries Dropped : 0
Total Neighbor Entries Aged-Out : 0
```

LOCAL-PORT	CHASSIS-ID	PORT-ID	PORT-DESC	TTL	SYS-NAME
1/1/8	08:00:09:5b:7e:2d	1/1/8	link to SW1	120	SW3
1/1/9	08:00:09:54:97:83	1/1/9	link to SW1	120	SW2

SW2

```
SW2# show lldp neighbor-info
```

```
LLDP Neighbor Information
=====
```

```
Total Neighbor Entries      : 2
Total Neighbor Entries Deleted : 0
Total Neighbor Entries Dropped : 0
```

Total Neighbor Entries Aged-Out : 0

LOCAL-PORT	CHASSIS-ID	PORT-ID	PORT-DESC	TTL	SYS-NAME
1/1/1	08:00:09:5b:7e:2d	1/1/9	link to SW2	120	SW3
1/1/9	08:00:09:d7:5f:0f	1/1/9	link to SW2	120	SW1

SW3

SW3# show lldp neighbor-info

LLDP Neighbor Information
=====

Total Neighbor Entries : 2
Total Neighbor Entries Deleted : 2
Total Neighbor Entries Dropped : 0
Total Neighbor Entries Aged-Out : 2

LOCAL-PORT	CHASSIS-ID	PORT-ID	PORT-DESC	TTL	SYS-NAME
1/1/8	08:00:09:d7:5f:0f	1/1/8	link to SW3	120	SW1
1/1/9	08:00:09:54:97:83	1/1/1	link to SW3	120	SW2

- Verify the IP connectivity with ping from SW1 to SW2 and SW3, from SW2 to SW1 and SW3, from SW3 to SW1 (some tests may be skipped):

SW1

```
SW1# ping 192.168.4.1
PING 192.168.4.1 (192.168.4.1) 100(128) bytes of data.
108 bytes from 192.168.4.1: icmp_seq=1 ttl=64 time=2.51 ms
108 bytes from 192.168.4.1: icmp_seq=2 ttl=64 time=2.73 ms
108 bytes from 192.168.4.1: icmp_seq=3 ttl=64 time=2.15 ms
108 bytes from 192.168.4.1: icmp_seq=4 ttl=64 time=2.17 ms
108 bytes from 192.168.4.1: icmp_seq=5 ttl=64 time=2.94 ms

--- 192.168.4.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4000ms
rtt min/avg/max/mdev = 2.148/2.499/2.941/0.309 ms
```

```
SW1# ping 192.168.4.11
PING 192.168.4.11 (192.168.4.11) 100(128) bytes of data.
108 bytes from 192.168.4.11: icmp_seq=1 ttl=64 time=3.26 ms
108 bytes from 192.168.4.11: icmp_seq=2 ttl=64 time=2.63 ms
108 bytes from 192.168.4.11: icmp_seq=3 ttl=64 time=2.40 ms
108 bytes from 192.168.4.11: icmp_seq=4 ttl=64 time=2.54 ms
108 bytes from 192.168.4.11: icmp_seq=5 ttl=64 time=2.09 ms

--- 192.168.4.11 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4003ms
rtt min/avg/max/mdev = 2.092/2.583/3.260/0.383 ms
```

```
SW1# ping 192.168.4.12
PING 192.168.4.12 (192.168.4.12) 100(128) bytes of data.
108 bytes from 192.168.4.12: icmp_seq=1 ttl=64 time=4.89 ms
108 bytes from 192.168.4.12: icmp_seq=2 ttl=64 time=8.16 ms
108 bytes from 192.168.4.12: icmp_seq=3 ttl=64 time=4.32 ms
108 bytes from 192.168.4.12: icmp_seq=4 ttl=64 time=3.95 ms
108 bytes from 192.168.4.12: icmp_seq=5 ttl=64 time=4.02 ms

--- 192.168.4.12 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 3.946/5.066/8.156/1.580 ms
```

SW2

```
SW2# ping 192.168.4.9
PING 192.168.4.9 (192.168.4.9) 100(128) bytes of data.
108 bytes from 192.168.4.9: icmp_seq=1 ttl=64 time=2.39 ms
108 bytes from 192.168.4.9: icmp_seq=2 ttl=64 time=2.46 ms
108 bytes from 192.168.4.9: icmp_seq=3 ttl=64 time=2.62 ms
108 bytes from 192.168.4.9: icmp_seq=4 ttl=64 time=2.15 ms
108 bytes from 192.168.4.9: icmp_seq=5 ttl=64 time=2.54 ms
```



```
--- 192.168.4.9 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 2.146/2.430/2.621/0.162 ms

SW2# ping 192.168.4.12
PING 192.168.4.12 (192.168.4.12) 100(128) bytes of data.
108 bytes from 192.168.4.12: icmp_seq=1 ttl=64 time=2.24 ms
108 bytes from 192.168.4.12: icmp_seq=2 ttl=64 time=2.50 ms
108 bytes from 192.168.4.12: icmp_seq=3 ttl=64 time=2.38 ms
108 bytes from 192.168.4.12: icmp_seq=4 ttl=64 time=1.85 ms
108 bytes from 192.168.4.12: icmp_seq=5 ttl=64 time=2.92 ms

--- 192.168.4.12 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 1.853/2.376/2.922/0.347 ms
```

SW3

```
SW3# ping 192.168.4.0
PING 192.168.4.0 (192.168.4.0) 100(128) bytes of data.
108 bytes from 192.168.4.0: icmp_seq=1 ttl=64 time=2.17 ms
108 bytes from 192.168.4.0: icmp_seq=2 ttl=64 time=2.37 ms
108 bytes from 192.168.4.0: icmp_seq=3 ttl=64 time=2.10 ms
108 bytes from 192.168.4.0: icmp_seq=4 ttl=64 time=2.24 ms
108 bytes from 192.168.4.0: icmp_seq=5 ttl=64 time=2.62 ms

--- 192.168.4.0 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 2.098/2.299/2.622/0.184 ms

SW3# ping 192.168.4.9
PING 192.168.4.9 (192.168.4.9) 100(128) bytes of data.
108 bytes from 192.168.4.9: icmp_seq=1 ttl=64 time=3.61 ms
108 bytes from 192.168.4.9: icmp_seq=2 ttl=64 time=2.94 ms
108 bytes from 192.168.4.9: icmp_seq=3 ttl=64 time=3.45 ms
108 bytes from 192.168.4.9: icmp_seq=4 ttl=64 time=3.93 ms
108 bytes from 192.168.4.9: icmp_seq=5 ttl=64 time=3.98 ms

--- 192.168.4.9 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4003ms
rtt min/avg/max/mdev = 2.937/3.579/3.980/0.377 ms
```

- Verify the IP connectivity between HostA and SW3 L0:

HostA

```
VPCS> ping 192.168.2.2

84 bytes from 192.168.2.2 icmp_seq=1 ttl=64 time=1.504 ms
84 bytes from 192.168.2.2 icmp_seq=2 ttl=64 time=1.149 ms
84 bytes from 192.168.2.2 icmp_seq=3 ttl=64 time=1.338 ms
84 bytes from 192.168.2.2 icmp_seq=4 ttl=64 time=1.511 ms
84 bytes from 192.168.2.2 icmp_seq=5 ttl=64 time=1.500 ms
```

Task 2 – Configure BGP routing

The objective of this task is to allow IP connectivity between HostA subnet and SW1 Loopback0. For this purpose, eBGP is used between SW1 and SW3. As there are 2 data paths, BGP weight will be used as selection criteria between the backup path (direct L3 link) and the main path which is through the intermediate L2 switch. Default VRF is used in this lab.

Step #1: Configure BGP

SW1 runs BGP Autonomous-System 65001, and Loopback0 is redistributed in BGP.

SW2 runs BGP Autonomous-System 65002, Loopback0 is redistributed in BGP as well as connected subnets that are redistributed based on ip prefix-list match.

SW1(config)#	SW3(config)#
<pre> router bgp 65001 bgp router-id 192.168.2.1 bgp fast-external-fallover neighbor 192.168.4.1 remote-as 65002 neighbor 192.168.4.1 weight 200 neighbor 192.168.4.12 remote-as 65002 neighbor 192.168.4.12 weight 300 neighbor 192.168.4.12 timers 10 30 address-family ipv4 unicast neighbor 192.168.4.1 activate neighbor 192.168.4.12 activate redistribute local loopback exit-address-family </pre>	<pre> ip prefix-list endpoint seq 10 permit 10.0.0.0/8 le 32 ! route-map connected-bgp permit seq 10 match ip address prefix-list endpoint ! router bgp 65002 bgp router-id 192.168.2.2 bgp fast-external-fallover neighbor 192.168.4.0 remote-as 65001 neighbor 192.168.4.0 weight 200 neighbor 192.168.4.9 remote-as 65001 neighbor 192.168.4.9 weight 300 neighbor 192.168.4.9 timers 10 30 address-family ipv4 unicast neighbor 192.168.4.0 activate neighbor 192.168.4.9 activate redistribute local loopback redistribute connected route-map connected- bgp exit-address-family </pre>

Note: the BGP timers for BGP sessions over the L2 intermediate node has been reduced here for minimizing recovery time during the lab. In production, default BGP timer values can be used as BFD would be used to protect session over intermediate L2 node.

Step #2: Verify BGP sessions

SW1

```
SW1# show bgp ipv4 unicast summary
```

```
VRF : default
```

```
GP Summary
```

```
-----
```

Local AS	: 65001	BGP Router Identifier	: 192.168.2.1
Peers	: 2	Log Neighbor Changes	: No
Cfg. Hold Time	: 180	Cfg. Keep Alive	: 60
Confederation Id	: 0		

Neighbor	Remote-AS	MsgRcvd	MsgSent	Up/Down	Time	State	AdminStatus
192.168.4.1	65002	74	74	00h:10m:03s		Established	Up
192.168.4.12	65002	74	76	00h:10m:03s		Established	Up

```
SW1# show bgp ipv4 unicast neighbors 192.168.4.12
```

```
Codes: ^ Inherited from peer-group
```

```
VRF : default
```

```
BGP Neighbor 192.168.4.12 (External)
```

```
Description      :
Peer-group       :
```

Remote Router Id	: 192.168.2.2	Local Router Id	: 192.168.2.1
Remote AS	: 65002	Local AS	: 65001
Remote Port	: 179	Local Port	: 44486
State	: Established	Admin Status	: Up
Conn. Established	: 1	Conn. Dropped	: 0
Passive	: No	Update-Source	:
Cfg. Hold Time	: 30	Cfg. Keep Alive	: 10
Neg. Hold Time	: 30	Neg. Keep Alive	: 10
Up/Down Time	: 00h:08m:08s	Connect-Retry Time	: 120
Local-AS Prepend	: No	Alt. Local-AS	: 0

```
BFD : Disabled
```

```

Password      :
Last Err Sent  : No Error
Last SubErr Sent : No Error
Last Err Rcvd  : No Error

```



```

Last SubErr Rcvd      : No Error

Graceful-Restart      : Enabled          Gr. Restart Time    : 120
Gr. Stalepath Time    : 300              Remove Private-AS   : No
TTL                   : 1                 Local Cluster-ID     : 
Weight                : 300              Fall-over           : No
Confederation-Peers   : No

Message statistics      Sent      Rcvd
-----
Open                   1          1
Notification           0          0
Updates                5          4
Keepalives             58         55
Route Refresh          0          0
Total                  64         60

Capability              Advertised   Received
-----
Route Refresh           Yes        Yes
Graceful Restart        Yes        Yes
Add-Path                No         No
Four Octet ASN          Yes        Yes
Address family IPv4 Unicast Yes        Yes
Address family IPv6 Unicast No         No
Address family L2VPN EVPN No         No

Address Family : IPv4 Unicast
-----

Rt. Reflect. Client    : No              Send Community      :
Allow-AS in            : 0                Advt. Interval      : 30
Max. Prefix            : 64000             Soft Reconfig In    :
Nexthop-Self           :                  Default-Originate    :
Cfg. Add-Path          :
Neg. Add-Path          : Disable

Routemap In            :
Routemap Out           :
ORF type               : Prefix-list
ORF capability         :

```

The BGP timer values displayed in the summary are the values for the global BGP timers (default 60s/180s).

Note: BFD is not yet enabled at this stage of the lab.

SW3

```

SW3# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----
Local AS                : 65002          BGP Router Identifier : 192.168.2.2
Peers                   : 2               Log Neighbor Changes  : No
Cfg. Hold Time          : 180             Cfg. Keep Alive       : 60
Confederation Id        : 0

Neighbor    Remote-AS  MsgRcvd  MsgSent  Up/Down Time  State      AdminStatus
192.168.4.0 65001       78       77       00h:10m:37s  Established Up
192.168.4.9 65001       81       78       00h:10m:37s  Established Up

SW3# show bgp ipv4 unicast neighbors 192.168.4.9
Codes: ^ Inherited from peer-group

VRF : default

BGP Neighbor 192.168.4.9 (External)
Description  :

```

```

Peer-group      :
Remote Router Id : 192.168.2.1
Remote AS       : 65001
Remote Port     : 44486
State           : Established
Conn. Established : 1
Passive         : No
Cfg. Hold Time  : 30
Neg. Hold Time  : 30
Up/Down Time    : 00h:11m:07s
Local-AS Prepend : No
BFD             : Disabled
Password        :

Last Err Sent   : No Error
Last SubErr Sent : No Error
Last Err Rcvd   : No Error
Last SubErr Rcvd : No Error

Graceful-Restart : Enabled
Gr. Stalepath Time : 300
TTL              : 1
Weight           : 300
Confederation-Peers : No

Local Router Id : 192.168.2.2
Local AS        : 65002
Local Port      : 179
Admin Status    : Up
Conn. Dropped   : 0
Update-Source   :
Cfg. Keep Alive : 10
Neg. Keep Alive : 10
Connect-Retry Time : 120
Alt. Local-AS   : 0

Message statistics      Sent      Rcvd
-----
Open                    1         1
Notification            0         0
Updates                 4         5
Keepalives              76        78
Route Refresh           0         0
Total                   81        84

Capability              Advertised      Received
-----
Route Refresh           Yes           Yes
Graceful Restart        Yes           Yes
Add-Path                No           No
Four Octet ASN          Yes           Yes
Address family IPv4 Unicast Yes           Yes
Address family IPv6 Unicast No            No
Address family L2VPN EVPN No            No

Address Family : IPv4 Unicast
-----

Rt. Reflect. Client : No
Allow-AS in         : 0
Max. Prefix         : 64000
Nextthop-Self       :
Cfg. Add-Path       :
Neg. Add-Path       : Disable

Routemap In         :
Routemap Out        :
ORF type            : Prefix-list
ORF capability       :

Send Community      :
Advt. Interval      : 30
Soft Reconfig In    :
Default-Originate   :

```

Step #3: Verify BGP routes and ip routing table

L0 of SW1 (192.168.2.1) should be learnt in SW3 behind two eBGP peers and only the route through SW2 should be preferred.

SW3

```

SW3# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

```

```
VRF : default
Local Router-ID 192.168.2.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*> 10.10.100.0/24	0.0.0.0	0	100	0	?
* e 192.168.2.1/32	192.168.4.0	0	100	200	65001 ?
*>e 192.168.2.1/32	192.168.4.9	0	100	300	65001 ?
*> 192.168.2.2/32	0.0.0.0	0	100	0	?

Total number of entries 4

Due to the higher weight of the BGP route learnt from the IP address of the peer circuit attached to the intermediate L2 switch, the path through this intermediate L2 switch is preferred in the routing table:

SW3

```
SW3# show ip route
```

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local

R - RIP, B - BGP, O - OSPF

Type Codes: E - External BGP, I - Internal BGP, V - VPN, EV - EVPN

IA - OSPF internal area, E1 - OSPF external type 1

E2 - OSPF external type 2

```
VRF: default
```

Prefix	Nexthop	Interface	VRF(egress)	Origin/ Type	Distance/ Metric	Age
10.10.100.0/24	-	vlan100	-	C	[0/0]	-
10.10.100.1/32	-	vlan100	-	L	[0/0]	-
192.168.2.1/32	192.168.4.9	vlan1105	-	B/E	[20/0]	01d:03h:52m
192.168.2.2/32	-	loopback0	-	L	[0/0]	-
192.168.4.0/31	-	1/1/8	-	C	[0/0]	-
192.168.4.1/32	-	1/1/8	-	L	[0/0]	-
192.168.4.8/29	-	vlan1105	-	C	[0/0]	-
192.168.4.12/32	-	vlan1105	-	L	[0/0]	-

Total Route Count : 8

Similarly, on SW1, the HostA subnet is preferred through the path of the intermediate L2 switch.

SW1

```
SW1# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,

i internal, e external S Stale, R Removed, a additional-paths

Origin codes: i - IGP, e - EGP, ? - incomplete

```
VRF : default
```

```
Local Router-ID 192.168.2.1
```

Network	Nexthop	Metric	LocPrf	Weight	Path
* e 10.10.100.0/24	192.168.4.1	0	100	200	65002 ?
*>e 10.10.100.0/24	192.168.4.12	0	100	300	65002 ?
*> 192.168.2.1/32	0.0.0.0	0	100	0	?
* e 192.168.2.2/32	192.168.4.1	0	100	200	65002 ?
*>e 192.168.2.2/32	192.168.4.12	0	100	300	65002 ?

Total number of entries 5

```
SW1# show ip route
```

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local

R - RIP, B - BGP, O - OSPF

Type Codes: E - External BGP, I - Internal BGP, V - VPN, EV - EVPN

IA - OSPF internal area, E1 - OSPF external type 1

E2 - OSPF external type 2

```
VRF: default
```

Prefix	Nexthop	Interface	VRF(egress)	Origin/ Type	Distance/ Metric	Age
10.10.100.0/24	192.168.4.12	vlan1105	-	B/E	[20/0]	01d:03h:56m
192.168.2.1/32	-	loopback0	-	L	[0/0]	-
192.168.2.2/32	192.168.4.12	vlan1105	-	B/E	[20/0]	01d:03h:56m
192.168.4.0/31	-	1/1/8	-	C	[0/0]	-
192.168.4.0/32	-	1/1/8	-	L	[0/0]	-
192.168.4.8/29	-	vlan1105	-	C	[0/0]	-
192.168.4.9/32	-	vlan1105	-	L	[0/0]	-
Total Route Count : 7						

Step #4: Verify IP connectivity between HostA and SW1 Loopback0

```

HostA
VPCS> ping 192.168.2.1

84 bytes from 192.168.2.1 icmp_seq=1 ttl=63 time=4.226 ms
84 bytes from 192.168.2.1 icmp_seq=2 ttl=63 time=4.424 ms
84 bytes from 192.168.2.1 icmp_seq=3 ttl=63 time=4.474 ms
84 bytes from 192.168.2.1 icmp_seq=4 ttl=63 time=3.991 ms
84 bytes from 192.168.2.1 icmp_seq=5 ttl=63 time=3.756 ms

```

Task 3 – Test resiliency without BFD

Let's perform a link failure of the main path, and check how long the traffic is interrupted before the backup path is established.

Step #1: Start pings from HostA to SW1 L0 until interrupted

```

HostA
VPCS> ping 192.168.2.1 -t

84 bytes from 192.168.2.1 icmp_seq=1 ttl=63 time=4.365 ms
84 bytes from 192.168.2.1 icmp_seq=2 ttl=63 time=4.259 ms
84 bytes from 192.168.2.1 icmp_seq=3 ttl=63 time=3.565 ms
84 bytes from 192.168.2.1 icmp_seq=4 ttl=63 time=4.187 ms
84 bytes from 192.168.2.1 icmp_seq=5 ttl=63 time=4.276 ms
84 bytes from 192.168.2.1 icmp_seq=6 ttl=63 time=3.852 ms
84 bytes from 192.168.2.1 icmp_seq=7 ttl=63 time=3.831 ms

```

Step #2: Modify VLAN ID on interface 1/1/9 of SW2 and check traffic interruption

```

SW2(config)#
SW2(config)# interface 1/1/9
SW2(config-if)# vlan access 1

```

Note: It can be a real scenario where the VLAN-ID on intermediate L2 switch is changed by mistake. There is no physical link transition and such event cannot be immediately detected by BGP without BFD.

Immediately after this change, you should see ICMP failures:

```

HostA
84 bytes from 192.168.2.1 icmp_seq=7 ttl=63 time=3.831 ms
84 bytes from 192.168.2.1 icmp_seq=8 ttl=63 time=5.068 ms
84 bytes from 192.168.2.1 icmp_seq=9 ttl=63 time=3.746 ms
84 bytes from 192.168.2.1 icmp_seq=10 ttl=63 time=4.111 ms
192.168.2.1 icmp_seq=11 timeout
192.168.2.1 icmp_seq=12 timeout
192.168.2.1 icmp_seq=13 timeout
192.168.2.1 icmp_seq=14 timeout

```

```

192.168.2.1 icmp_seq=15 timeout
192.168.2.1 icmp_seq=16 timeout
192.168.2.1 icmp_seq=17 timeout
192.168.2.1 icmp_seq=18 timeout
192.168.2.1 icmp_seq=19 timeout
192.168.2.1 icmp_seq=20 timeout
...
192.168.2.1 icmp_seq=83 timeout
192.168.2.1 icmp_seq=84 timeout
84 bytes from 192.168.2.1 icmp_seq=85 ttl=63 time=2.043 ms
84 bytes from 192.168.2.1 icmp_seq=86 ttl=63 time=2.481 ms
84 bytes from 192.168.2.1 icmp_seq=87 ttl=63 time=2.715 ms
84 bytes from 192.168.2.1 icmp_seq=88 ttl=63 time=6.122 ms
84 bytes from 192.168.2.1 icmp_seq=89 ttl=63 time=2.920 ms
84 bytes from 192.168.2.1 icmp_seq=90 ttl=63 time=3.013 ms

```

It should take about the BGP hold-time duration for the alternate path to become selected by BGP.

Conclusion: without BFD, this incident on the intermediate L2 switch creates a significant network outage.

Step #3: restore VLAN 1105 on interface 1/1/9 of SW2

Check that routing is back to nominal state.

Step #4 (optional): test bgp fast-external-fallover

Optionally, you can compare routing responsiveness in case of direct L3 circuit failure and test the effect of BGP fast-external-fallover, by inverting main and backup path. For this purpose, you can modify the weight of the backup-path to become the active path.

SW1(config)#	SW3(config)#
<pre> router bgp 65001 bgp router-id 192.168.2.1 bgp fast-external-fallover neighbor 192.168.4.1 remote-as 65002 neighbor 192.168.4.1 weight 400 neighbor 192.168.4.12 remote-as 65002 neighbor 192.168.4.12 weight 300 neighbor 192.168.4.12 timers 10 30 </pre>	<pre> router bgp 65002 bgp router-id 192.168.2.2 bgp fast-external-fallover neighbor 192.168.4.0 remote-as 65001 neighbor 192.168.4.0 weight 400 neighbor 192.168.4.9 remote-as 65001 neighbor 192.168.4.9 weight 300 neighbor 192.168.4.9 timers 10 30 </pre>

Proceed with “clear bgp **” on both SW1 and SW3 for the new weight value to take effect.

You can repeat the ICMP test and shutdown interface 1/1/8 **on both** SW1 and SW3. This interface shutdown must be executed on both SW1 and SW3 as **simultaneously as possible** (reminder: there is no link state follow-up if peer link is shutdown on the CX Simulator).

The duration for traffic rerouting should be minimum: below 1 or 2 seconds ICMP traffic interruption depending on how simultaneous is the interface shutdown executed on both SW1 and SW3.

Note: BFD is not needed for direct L3 circuit between two eBGP peers.

Restore previous weight value and restore interface 1/1/8 on both SW1 and SW3 if you ran this optional test.

Task 4 – Configure BFD for the eBGP session established over the L2 path

- BFD must be enabled globally.
- BFD is configured globally and not per interface.
The BGP network-service will create the BFD session on the associated interface of the BGP peering.
- BFD must be configured on both ends (SW1 and SW3). If only one end is configured, the BFD state will stay down (instead of going through init, then up state).

- Echo loopback function is not available on CX Simulator.
- Proposed timers are the minimum value to obtain a reasonable failure detection time (1~2 seconds).

Step #1: Configure BFD globally

SW1(config)#	SW3(config)#
bfd	bfd
bfd min-receive-interval 500	bfd min-receive-interval 500
bfd min-transmit-interval 500	bfd min-transmit-interval 500
bfd detect-multiplier 3	bfd detect-multiplier 3

Step #2: Configure BFD for eBGP session

SW1(config)#	SW3(config)#
router bgp 65001	router bgp 65002
bgp router-id 192.168.2.1	bgp router-id 192.168.2.2
bgp fast-external-fallover	bgp fast-external-fallover
neighbor 192.168.4.1 remote-as 65002	neighbor 192.168.4.0 remote-as 65001
neighbor 192.168.4.1 weight 200	neighbor 192.168.4.0 weight 200
neighbor 192.168.4.12 remote-as 65002	neighbor 192.168.4.9 remote-as 65001
neighbor 192.168.4.12 weight 300	neighbor 192.168.4.9 weight 300
neighbor 192.168.4.12 timers 10 30	neighbor 192.168.4.9 timers 10 30
neighbor 192.168.4.12 fall-over bfd	neighbor 192.168.4.9 fall-over bfd

No need to clear the BGP session, the BFD session is created immediately.

Step #3: Verify BFD

Verify that BFD is set on the BGD session:

SW1			
SW1# show bgp ipv4 unicast neighbors 192.168.4.12			
Codes: ^ Inherited from peer-group			
VRF : default			
BGP Neighbor 192.168.4.12 (External)			
Description :			
Peer-group :			
Remote Router Id	: 192.168.2.2	Local Router Id	: 192.168.2.1
Remote AS	: 65002	Local AS	: 65001
Remote Port	: 49390	Local Port	: 179
State	: Established	Admin Status	: Up
Conn. Established	: 2	Conn. Dropped	: 1
Passive	: No	Update-Source	:
Cfg. Hold Time	: 30	Cfg. Keep Alive	: 10
Neg. Hold Time	: 30	Neg. Keep Alive	: 10
Up/Down Time	: 00h:20m:30s	Connect-Retry Time	: 120
Local-AS Prepend	: No	Alt. Local-AS	: 0
BFD	: Enabled		
Password :			
Last Err Sent	: No Error		
Last SubErr Sent	: No Error		
Last Err Rcvd	: No Error		
Last SubErr Rcvd	: No Error		
Graceful-Restart	: Enabled	Gr. Restart Time	: 120
Gr. Stalepath Time	: 300	Remove Private-AS	: No
TTL	: 1	Local Cluster-ID	:
Weight	: 300	Fall-over	: No
Confederation-Peers	: No		
Message statistics	Sent	Rcvd	


```

-----
Open                2          2
Notification        0          0
Updates             7          6
Keepalives          148        149
Route Refresh       0          0
Total               157        157

Capability
-----
Route Refresh       Yes
Graceful Restart    Yes
Add-Path            No
Four Octet ASN      Yes
Address family IPv4 Unicast Yes
Address family IPv6 Unicast No
Address family L2VPN EVPN  No

Address Family : IPv4 Unicast
-----

Rt. Reflect. Client : No
Allow-AS in         : 0
Max. Prefix         : 64000
Nexthop-Self        :
Cfg. Add-Path       :
Neg. Add-Path       : Disable

Routemap In         :
Routemap Out        :
ORF type            : Prefix-list
ORF capability       :

Send Community      :
Advt. Interval      : 30
Soft Reconfig In    :
Default-Originate   :

```

SW3

```

SW3# sh bgp ipv4 unicast neighbors 192.168.4.9
Codes: ^ Inherited from peer-group

```

VRF : default

BGP Neighbor 192.168.4.9 (External)

```

Description          :

Peer-group           :

Remote Router Id     : 192.168.2.1    Local Router Id      : 192.168.2.2
Remote AS            : 65001          Local AS             : 65002
Remote Port          : 179            Local Port           : 49390
State                : Established    Admin Status         : Up
Conn. Established    : 1              Conn. Dropped        : 0
Passive              : No             Update-Source         :
Cfg. Hold Time       : 30             Cfg. Keep Alive      : 10
Neg. Hold Time       : 30             Neg. Keep Alive       : 10
Up/Down Time         : 00h:22m:42s    Connect-Retry Time   : 120
Local-AS Prepend     : No             Alt. Local-AS        : 0
BFD                  : Enabled
Password             :

Last Err Sent        : No Error
Last SubErr Sent     : No Error
Last Err Rcvd        : No Error
Last SubErr Rcvd     : No Error

Graceful-Restart      : Enabled        Gr. Restart Time     : 120
Gr. Stalepath Time   : 300            Remove Private-AS    : No
TTL                  : 1              Local Cluster-ID     :
Weight               : 300            Fall-over            : No
Confederation-Peers  : No

```

```

Message statistics    Sent      Rcvd
-----
Open                  1          1
Notification          0          0
Updates               4          2

```

Keepalives	156	156	
Route Refresh	0	0	
Total	161	159	
Capability		Advertised	Received
-----		-----	-----
Route Refresh	Yes	Yes	Yes
Graceful Restart	Yes	Yes	Yes
Add-Path	No	No	No
Four Octet ASN	Yes	Yes	Yes
Address family IPv4 Unicast	Yes	Yes	Yes
Address family IPv6 Unicast	No	No	No
Address family L2VPN EVPN	No	No	No
Address Family : IPv4 Unicast			

Rt. Reflect. Client : No		Send Community :	
Allow-AS in : 0		Advt. Interval : 30	
Max. Prefix : 64000		Soft Reconfig In :	
Nexthop-Self :		Default-Originate :	
Cfg. Add-Path :			
Neg. Add-Path : Disable			
Routemap In :			
Routemap Out :			
ORF type : Prefix-list			
ORF capability :			

Then verify BFD with various show bfd commands:

SW1

SW1# show bfd summary

Sessions

Admin	Down	Down	Init	Up	Total
	0	0	0	1	1

Total protocols: 0

SW1# show bfd

Admin status: enabled

Echo source IP: N/A

Statistics:

Total number of control packets transmitted: 101135

Total number of control packets received: 101088

Total number of control packets dropped: 0

Session	Interface	VRF	Source IP	Destination IP	Echo	State	Protocol
1	vlan1105	default	192.168.4.9	192.168.4.12	N/A	up	bgp

SW1# show bfd session 1

BFD session information - Session 1

VRF: default

Min Tx interval (msec): 500

Min Rx interval (msec): 500

Min echo Rx interval (msec): 500

Detect multiplier: 3

Protocol(s): bgp

Local discriminator: 4310

Remote discriminator: 13211

Echo: N/A

```
Local diagnostic: no_diagnostic
Remote diagnostic: no_diagnostic
State flaps: 0
Interface Source IP      Destination IP      State      Pkt Rx      Pkt Tx      Pkt drop
-----
vlan1105 192.168.4.9        192.168.4.12      up         102984      103030      0
```

Please note the discriminator field. These values will be used in a second BFD lab while tracing the BFD control packets.

Similarly, on SW3:

SW3

```
SW3# show bfd summary
```

Sessions

```
-----
Admin Down    Down    Init    Up    Total
-----
              0        0        0        1        1
-----
```

```
Total protocols: 0
```

```
SW3# show bfd
```

```
Admin status: enabled
```

```
Echo source IP: N/A
```

```
Statistics:
```

```
Total number of control packets transmitted: 102595
```

```
Total number of control packets received: 102592
```

```
Total number of control packets dropped: 0
```

```
Session Interface VRF      Source IP      Destination IP      Echo      State      Protocol
-----
1      vlan1105  default      192.168.4.12      192.168.4.9        N/A      up         bgp
```

```
SW3# show bfd session 1
```

```
BFD session information - Session 1
```

```
VRF: default
```

```
Min Tx interval (msec): 500
```

```
Min Rx interval (msec): 500
```

```
Min echo Rx interval (msec): 500
```

```
Detect multiplier: 3
```

```
Protocol(s): bgp
```

```
Local discriminator: 13211
```

```
Remote discriminator: 4310
```

```
Echo: N/A
```

```
Local diagnostic: no_diagnostic
```

```
Remote diagnostic: no_diagnostic
```

```
State flaps: 0
```

```
Interface Source IP      Destination IP      State      Pkt Rx      Pkt Tx      Pkt drop
-----
vlan1105 192.168.4.12        192.168.4.9        up         104557      104558      0
```

Note that the discriminator values reported by SW3 are the mirrored values of SW1.

Task 5 – Test resiliency with BFD

Step #1: Start pings from HostA to SW1 L0 until interrupted

HostA

```
VPCS> ping 192.168.2.1 -t
```

```
84 bytes from 192.168.2.1 icmp_seq=1 ttl=63 time=4.365 ms
```

```
84 bytes from 192.168.2.1 icmp_seq=2 ttl=63 time=4.259 ms
```

```
84 bytes from 192.168.2.1 icmp_seq=3 ttl=63 time=3.565 ms
84 bytes from 192.168.2.1 icmp_seq=4 ttl=63 time=4.187 ms
84 bytes from 192.168.2.1 icmp_seq=5 ttl=63 time=4.276 ms
84 bytes from 192.168.2.1 icmp_seq=6 ttl=63 time=3.852 ms
84 bytes from 192.168.2.1 icmp_seq=7 ttl=63 time=3.831 ms
```

Step #2: Modify VLAN ID on interface 1/1/9 of SW2 and check traffic interruption

SW2(config)#

```
SW2(config)# interface 1/1/9
SW2(config-if)# vlan access 1
```

Immediately after this change, you should see almost no ICMP failure (no more than 1):

HostA

```
84 bytes from 192.168.2.1 icmp_seq=8 ttl=63 time=3.936 ms
84 bytes from 192.168.2.1 icmp_seq=9 ttl=63 time=3.995 ms
84 bytes from 192.168.2.1 icmp_seq=10 ttl=63 time=2.810 ms
84 bytes from 192.168.2.1 icmp_seq=11 ttl=63 time=3.605 ms
84 bytes from 192.168.2.1 icmp_seq=12 ttl=63 time=3.430 ms
84 bytes from 192.168.2.1 icmp_seq=13 ttl=63 time=4.012 ms
84 bytes from 192.168.2.1 icmp_seq=14 ttl=63 time=7.406 ms
84 bytes from 192.168.2.1 icmp_seq=15 ttl=63 time=4.186 ms
84 bytes from 192.168.2.1 icmp_seq=16 ttl=63 time=3.866 ms
84 bytes from 192.168.2.1 icmp_seq=17 ttl=63 time=4.532 ms
192.168.2.1 icmp_seq=18 timeout
84 bytes from 192.168.2.1 icmp_seq=19 ttl=63 time=4.061 ms
84 bytes from 192.168.2.1 icmp_seq=20 ttl=63 time=2.446 ms
84 bytes from 192.168.2.1 icmp_seq=21 ttl=63 time=2.482 ms
84 bytes from 192.168.2.1 icmp_seq=22 ttl=63 time=3.593 ms
84 bytes from 192.168.2.1 icmp_seq=23 ttl=63 time=3.058 ms
84 bytes from 192.168.2.1 icmp_seq=24 ttl=63 time=3.174 ms
```

It should take 1.5 second (3*500ms) to detect BGP BFD failure which tears down the BGP session to the "connect" state.

Step #3: Check status of BFD session and BGP session. Check routes.

Look at the status of BFD session as well as the effect on the BGP session.

SW1

SW1# show bfd session 1

BFD session information - Session 1

VRF: default

Min Tx interval (msec): 1000

Min Rx interval (msec): 500

Min echo Rx interval (msec): 500

Detect multiplier: 3

Protocol(s): bgp

Local discriminator: 4310

Remote discriminator: 0

Echo: N/A

Local diagnostic: control_detection_time_expired

Remote diagnostic: no_diagnostic

State flaps: 1

Interface	Source IP	Destination IP	State	Pkt Rx	Pkt Tx	Pkt drop
vlan1105	192.168.4.9	192.168.4.12	down	107331	107396	0

SW1# show bgp ipv4 unicast summary

VRF : default

BGP Summary

Local AS	: 65001	BGP Router Identifier	: 192.168.2.1
Peers	: 2	Log Neighbor Changes	: No
Cfg. Hold Time	: 180	Cfg. Keep Alive	: 60
Confederation Id	: 0		

Neighbor	Remote-AS	MsgRcvd	MsgSent	Up/Down Time	State	AdminStatus
192.168.4.1	65002	92	92	01h:05m:40s	Established	Up
192.168.4.12	65002	460	460	00h:01m:12s	Connect	Up

SW3

SW3# show bfd session 1

BFD session information - Session 1

VRF: default

Min Tx interval (msec): 1000

Min Rx interval (msec): 500

Min echo Rx interval (msec): 500

Detect multiplier: 3

Protocol(s): bgp

Local discriminator: 13211

Remote discriminator: 0

Echo: N/A

Local diagnostic: control_detection_time_expired

Remote diagnostic: no_diagnostic

State flaps: 1

Interface	Source IP	Destination IP	State	Pkt Rx	Pkt Tx	Pkt drop
vlan1105	192.168.4.12	192.168.4.9	down	107329	107366	0

SW3# show bgp ipv4 unicast summary

VRF : default

BGP Summary

Local AS	: 65002	BGP Router Identifier	: 192.168.2.2
Peers	: 2	Log Neighbor Changes	: No
Cfg. Hold Time	: 180	Cfg. Keep Alive	: 60
Confederation Id	: 0		

Neighbor	Remote-AS	MsgRcvd	MsgSent	Up/Down Time	State	AdminStatus
192.168.4.0	65001	81	79	01h:05m:15s	Established	Up
192.168.4.9	65001	447	449	00h:00m:47s	Connect	Up

Consequently, the alternate BGP routes already present in the BGP table are selected as best and populate the ip routing table:

SW1

SW1# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local

R - RIP, B - BGP, O - OSPF

Type Codes: E - External BGP, I - Internal BGP, V - VPN, EV - EVPN

IA - OSPF internal area, E1 - OSPF external type 1

E2 - OSPF external type 2

VRF: default

Prefix	Nexthop	Interface	VRF(egress)	Origin/ Type	Distance/ Metric	Age
10.10.100.0/24	192.168.4.1	1/1/8	-	B/E	[20/0]	00h:01m:20s
192.168.2.1/32	-	loopback0	-	L	[0/0]	-
192.168.2.2/32	192.168.4.1	1/1/8	-	B/E	[20/0]	00h:01m:20s
192.168.4.0/31	-	1/1/8	-	C	[0/0]	-
192.168.4.0/32	-	1/1/8	-	L	[0/0]	-
192.168.4.8/29	-	vlan1105	-	C	[0/0]	-
192.168.4.9/32	-	vlan1105	-	L	[0/0]	-

Total Route Count : 7

SW3

SW3# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local

R - RIP, B - BGP, O - OSPF

Type Codes: E - External BGP, I - Internal BGP, V - VPN, EV - EVPN

IA - OSPF internal area, E1 - OSPF external type 1

E2 - OSPF external type 2

VRF: default

Prefix	NextHop	Interface	VRF(egress)	Origin/ Type	Distance/ Metric	Age
10.10.100.0/24	-	vlan100	-	C	[0/0]	-
10.10.100.1/32	-	vlan100	-	L	[0/0]	-
192.168.2.1/32	192.168.4.0	1/1/8	-	B/E	[20/0]	00h:00m:59s
192.168.2.2/32	-	loopback0	-	L	[0/0]	-
192.168.4.0/31	-	1/1/8	-	C	[0/0]	-
192.168.4.1/32	-	1/1/8	-	L	[0/0]	-
192.168.4.8/29	-	vlan1105	-	C	[0/0]	-
192.168.4.12/32	-	vlan1105	-	L	[0/0]	-

Total Route Count : 8

Conclusion: with BFD, this incident on the intermediate L2 switch has a minimum impact on the network.

This is the end of this lab.

Appendix – Reference Configurations

If you face issues during your lab, you can verify your configuration with the configuration extract listed in this section.

SW1

```
hostname SW1
!
bfd
bfd min-receive-interval 500
bfd min-transmit-interval 500
bfd detect-multiplier 3
!
vlan 1
vlan 1105
    description transit interco VLAN
interface mgmt
    no shutdown
    ip dhcp
interface 1/1/8
    no shutdown
    description link to SW3
    ip address 192.168.4.0/31
interface 1/1/9
    no shutdown
    description link to SW2
    no routing
    vlan access 1105
interface loopback 0
    ip address 192.168.2.1/32
interface vlan 1105
    ip address 192.168.4.9/29
!
router bgp 65001
    bgp router-id 192.168.2.1
    bgp fast-external-fallover
    neighbor 192.168.4.1 remote-as 65002
    neighbor 192.168.4.1 weight 200
    neighbor 192.168.4.12 remote-as 65002
    neighbor 192.168.4.12 weight 300
    neighbor 192.168.4.12 timers 10 30
    neighbor 192.168.4.12 fall-over bfd
    address-family ipv4 unicast
        neighbor 192.168.4.1 activate
        neighbor 192.168.4.12 activate
        redistribute local loopback
    exit-address-family
```

SW2

```
hostname SW2
!
vlan 1
vlan 1105
    description transit interco VLAN
interface mgmt
    no shutdown
    ip dhcp
interface 1/1/1
    no shutdown
    description link to SW3
    no routing
    vlan access 1105
interface 1/1/9
    no shutdown
    description link to SW1
    no routing
    vlan access 1105
interface vlan 1105
    ip address 192.168.4.11/29
```

SW3

```
hostname SW3
!
bfd
bfd min-receive-interval 500
bfd min-transmit-interval 500
bfd detect-multiplier 3
!
vlan 1,100
vlan 1105
    description transit interco VLAN
interface mgmt
    no shutdown
    ip dhcp
interface 1/1/1
    no shutdown
    description link to HostA
    no routing
    vlan access 100
interface 1/1/8
    no shutdown
    description link to SW1
    ip address 192.168.4.1/31
interface 1/1/9
    no shutdown
    description link to SW2
    no routing
    vlan access 1105
interface loopback 0
    ip address 192.168.2.2/32
interface vlan 100
    ip address 10.10.100.1/24
interface vlan 1105
    ip address 192.168.4.12/29
!
ip prefix-list endpoint seq 10 permit 10.0.0.0/8 le 32
!
route-map connected-bgp permit seq 10
    match ip address prefix-list endpoint
!
router bgp 65002
    bgp router-id 192.168.2.2
    bgp fast-external-fallover
    neighbor 192.168.4.0 remote-as 65001
    neighbor 192.168.4.0 weight 200
    neighbor 192.168.4.9 remote-as 65001
    neighbor 192.168.4.9 weight 300
    neighbor 192.168.4.9 timers 10 30
    neighbor 192.168.4.9 fall-over bfd
    address-family ipv4 unicast
        neighbor 192.168.4.0 activate
        neighbor 192.168.4.9 activate
        redistribute local loopback
        redistribute connected route-map connected-bgp
    exit-address-family
```

